

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Methods and Systems for Implementing Dynamic
Properties on Objects that Support Only Static
Properties**

Inventor(s):

Daniel J. Miller

David Maymudes

ATTORNEY'S DOCKET NO. ms1-638us

1 **TECHNICAL FIELD**

2
3 This invention generally relates to processing media content and, more
4 particularly, to a system and related interfaces facilitating the processing of media
5 content.

6 **BACKGROUND**

7
8 Recent advances in computing power and related technology have fostered
9 the development of a new generation of powerful software applications. Gaming
10 applications, communications applications, and multimedia applications have
11 particularly benefited from increased processing power and clocking speeds.
12 Indeed, once the province of dedicated, specialty workstations, many personal
13 computing systems now have the capacity to receive, process and render
14 multimedia objects (e.g., audio and video content). While the ability to display
15 (receive, process and render) multimedia content has been around for a while, the
16 ability for a standard computing system to support true multimedia editing
17 applications is relatively new.

18 In an effort to satisfy this need, Microsoft Corporation introduced an
19 innovative development system supporting advanced user-defined multimedia
20 editing functions. An example of this architecture is presented in US Patent No.
21 5,913, 038 issued to Griffiths and commonly owned by the assignee of the present
22 invention, the disclosure of which is expressly incorporated herein by reference.

23 In the '038 patent, Griffiths introduced the an application program interface
24 which, when exposed to higher-level development applications, enables a user to
25 graphically construct a multimedia processing project by piecing together a

1 collection of "filters" exposed by the interface. The interface described therein is
2 referred to as a filter graph manager. The filter graph manager controls the data
3 structure of the filter graph and the way data moves through the filter graph. The
4 filter graph manager provides a set of component object model (COM) interfaces
5 for communication between a filter graph and its application. Filters of a filter
6 graph architecture are preferably implemented as COM objects, each
7 implementing one or more interfaces, each of which contains a predefined set of
8 functions, called methods. Methods are called by an application program or other
9 component objects in order to communicate with the object exposing the interface.
10 The application program can also call methods or interfaces exposed by the filter
11 graph manager object.

12 Filter graphs work with data representing a variety of media (or non-media)
13 data types, each type characterized by a data stream that is processed by the filter
14 components comprising the filter graph. A filter positioned closer to the source of
15 the data is referred to as an upstream filter, while those further down the
16 processing chain is referred to as a downstream filter. For each data stream that
17 the filter handles it exposes at least one virtual pin (i.e., distinguished from a
18 physical pin such as one might find on an integrated circuit). A virtual pin can be
19 implemented as a COM object that represents a point of connection for a
20 unidirectional data stream on a filter. Input pins represent inputs and accept data
21 into the filter, while output pins represent outputs and provide data to other filters.
22 Each of the filters include at least one memory buffer, wherein communication of
23 the media stream between filters is often accomplished by a series of "copy"
24 operations from one filter to another.

As introduced in Griffiths, a filter graph has three different types of filters: source filters, transform filters, and rendering filters. A source filter is used to load data from some source; a transform filter processes and passes data; and a rendering filter renders data to a hardware device or other locations (e.g., saved to a file, etc.). An example of a filter graph for a simplistic media rendering process is presented with reference to Fig. 1.

Fig. 1 graphically illustrates an example filter graph for rendering media content. As shown, the filter graph 100 is comprised of a plurality of filters 102-114, which read, process (transform) and render media content from a selected source file. As shown, the filter graph includes each of the types of filters described above, interconnected in a linear fashion.

Products utilizing the filter graph have been well received in the market as it has opened the door to multimedia editing using otherwise standard computing systems. It is to be appreciated, however, that the construction and implementation of the filter graphs are computationally intensive and expensive in terms of memory usage. Even the most simple of filter graphs requires and abundance of memory to facilitate the copy operations required to move data between filters. Complex filter graphs can become unwieldy, due in part to the linear nature of prior art filter graph architecture. Moreover, it is to be appreciated that the filter graphs themselves consume memory resources, thereby compounding the issue introduced above.

Thus, what is required is a filter graph architecture which reduces the computational and memory resources required to support even the most complex of multimedia projects. Indeed, what is required is a dynamically reconfigurable

multimedia editing system and related methods, unencumbered by the limitations described above. Just such a system and methods are disclosed below.

SUMMARY

Methods and systems of simulating dynamic properties on computer-implemented objects that do not support dynamic properties are described. In one embodiment, one or more first objects that do not support dynamic properties are provided. One or more second programmable objects are provided and are programmed to effect property value changes on the objects that do not support dynamic properties. The programmable objects can be programmed using data structures that, in one embodiment, comprise an array of one or more sets of data structures. Each data structure set is associated with a property whose value is desired to be changed. The data structure set can define a new property value, a time at which the property value is to be changed, and how to effect the property value change. The programmable object(s) is pre-programmed with the data structures and knows when to call the first objects so that they can change their properties. In one embodiment, the programmable objects are employed in the context of multi-media project editing software that permits a user to build a multi-media project using multiple different digital source streams.

BRIEF DESCRIPTION OF THE DRAWINGS

The same reference numbers are used throughout the figures to reference like components and features.

Fig. 2 is a block diagram of a computing system incorporating the teachings of the described embodiment.

Fig. 3 is a block diagram of an example software architecture incorporating the teachings of the described embodiment.

Fig. 4 is a graphical illustration of an example software-enabled matrix switch, according to an exemplary embodiment.

Fig. 5 is a graphical representation of a data structure comprising a programming grid to selectively couple one or more of a scalable plurality of input pins to a scalable plurality of output pins of the matrix switch filter, in accordance with one aspect of the described embodiment.

Fig. 6 is a graphical illustration denoting shared buffer memory between filters, according to one aspect of the described embodiment.

Fig. 7 is a flow chart of an example method for generating a filter graph, in accordance with one aspect of the described embodiment.

Fig. 8 is a flow chart of an example method for negotiating buffer requirements between at least two adjacent filters, according to one aspect of the described embodiment.

Fig. 9 graphically illustrates an overview of a process that takes a user-defined editing project and composites a data structure that can be used to program the matrix switch.

Fig. 10 graphically illustrates the project of Fig. 9 in greater detail.

Fig. 11 shows an exemplary matrix switch dynamically generated in support of the project developed in Figs. 9 and 10, according to one described embodiment.

Fig. 12 illustrates a graphic representation of an exemplary data structure that represents the project of Fig. 10, according to one described embodiment.

Figs. 13-18 graphically illustrate various states of a matrix switch programming grid at select points in processing the project of Figs. 9 and 10 through the matrix switch, in accordance with one described embodiment.

Fig. 19 is a flow chart of an example method for processing media content, in accordance with one described embodiment.

Fig. 20 illustrates an example project with a transition and an effect, in accordance with one described embodiment.

Fig. 21 shows an exemplary data structure in the form of a hierarchical tree that represents the project of Fig. 20.

Figs. 22 and 23 graphically illustrate an example matrix switch programming grid associated with the project of Fig. 20 at select points in time, according to one described embodiment.

Fig. 24 shows an example matrix switch dynamically generated and configured as the grid of Figs. 22 and 23 was being processed, in accordance with one described embodiment.

Fig. 25 shows an exemplary project in accordance with one described embodiment.

Fig. 26 graphically illustrates an example audio editing project, according to one described embodiment.

Fig. 39 is a flow diagram that describes steps in a method in accordance with one described embodiment.

1 **DETAILED DESCRIPTION**

2 **Related Applications**

3 This application is related to the following commonly-filed U.S. Patent
4 Applications, all of which are commonly assigned to Microsoft Corp., the
5 disclosures of which are incorporated by reference herein:
6

- 7
- Application Serial No. _____, entitled "An Interface and Related Methods for Reducing Source Accesses in a Development System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-643US;
 - Application Serial No. _____, entitled "A System and Related Interfaces Supporting the Processing of Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-629US;
 - Application Serial No. _____, entitled "A System and Related Methods for Reducing Source Filter Invocation in a Development Project", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-631US;
 - Application Serial No. _____, entitled "A System and Related Methods for Reducing Memory Requirements of a Media Processing System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-632US;
 - Application Serial No. _____, entitled "A System and Related Methods for Reducing the Instances of Source Files in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-633US;
 - Application Serial No. _____, entitled "An Interface and Related Methods for Dynamically Generating a Filter Graph in a Development System", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-634US;
 - Application Serial No. _____, entitled "A System and Related Methods for Processing Audio Content in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-639US;
 - Application Serial No. _____, entitled "A System and Methods for Generating an Managing Filter Strings in a Filter Graph", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-642US;
- 24
25

- Application Serial No. _____, entitled "Methods and Systems for Processing Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-640US;
- Application Serial No. _____, entitled "Systems for Managing Multiple Inputs and Methods and Systems for Processing Media Content ", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-635US;
- Application Serial No. _____, entitled "Methods and Systems for Implementing Dynamic Properties on Objects that Support Only Static Properties", naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-638US;
- Application Serial No. _____, entitled "Methods and Systems for Efficiently Processing Compressed and Uncompressed Media Content", naming Daniel J. Miller and Eric H. Rudolph as inventors, and bearing attorney docket number MS1-630US;
- Application Serial No. _____, entitled "Methods and Systems for Effecting Video Transitions Represented By Bitmaps", naming Daniel J. Miller and David Maymudes as inventors, and bearing attorney docket number MS1-637US;
- Application Serial No. _____, entitled "Methods and Systems for Mixing Digital Audio Signals", naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-636US; and
- Application Serial No. _____, entitled "Methods and Systems for Processing Multi-media Editing Projects", naming Eric H. Rudolph as inventor, and bearing attorney docket number MS1-641US.

Various described embodiments concern an application program interface associated with a development system. According to one example implementation, the interface is exposed to a media processing application to enable a user to dynamically generate complex media processing tasks, e.g., editing projects. In the discussion herein, aspects of the invention are developed within the general context of computer-executable instructions, such as program modules, being executed by one or more conventional computers. Generally, program modules include routines, programs, objects, components, data structures,

etc. that perform particular tasks or implement particular abstract data types. Moreover, those skilled in the art will appreciate that the invention may be practiced with other computer system configurations, including hand-held devices, personal digital assistants, multiprocessor systems, microprocessor-based or programmable consumer electronics, network PCs, minicomputers, mainframe computers, and the like. In a distributed computer environment, program modules may be located in both local and remote memory storage devices. It is noted, however, that modification to the architecture and methods described herein may well be made without deviating from spirit and scope of the present invention. Moreover, although developed within the context of a media processing system paradigm, those skilled in the art will appreciate, from the discussion to follow, that the application program interface may well be applied to other development system implementations. Thus, the media processing system described below is but one illustrative implementation of a broader inventive concept.

Example System Architecture

Fig. 2 illustrates an example of a suitable computing environment 200 on which the system and related methods for processing media content may be implemented.

It is to be appreciated that computing environment 200 is only one example of a suitable computing environment and is not intended to suggest any limitation as to the scope of use or functionality of the media processing system. Neither should the computing environment 200 be interpreted as having any dependency or requirement relating to any one or combination of components illustrated in the exemplary computing environment 200.

The media processing system is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well known computing systems, environments, and/or configurations that may be suitable for use with the media processing system include, but are not limited to, personal computers, server computers, thin clients, thick clients, handheld or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputers, mainframe computers, distributed computing environments that include any of the above systems or devices, and the like.

In certain implementations, the system and related methods for processing media content may well be described in the general context of computer-executable instructions, such as program modules, being executed by a computer. Generally, program modules include routines, programs, objects, components, data structures, etc. that perform particular tasks or implement particular abstract data types. The media processing system may also be practiced in distributed computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed computing environment, program modules may be located in both local and remote computer storage media including memory storage devices.

In accordance with the illustrated example embodiment of Fig. 2 computing system 200 is shown comprising one or more processors or processing units 202, a system memory 204, and a bus 206 that couples various system components including the system memory 204 to the processor 202.

Bus 206 is intended to represent one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an

1 accelerated graphics port, and a processor or local bus using any of a variety of
2 bus architectures. By way of example, and not limitation, such architectures
3 include Industry Standard Architecture (ISA) bus, Micro Channel Architecture
4 (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association
5 (VESA) local bus, and Peripheral Component Interconnects (PCI) buss also
6 known as Mezzanine bus.

7 Computer 200 typically includes a variety of computer readable media.
8 Such media may be any available media that is locally and/or remotely accessible
9 by computer 200, and it includes both volatile and non-volatile media, removable
10 and non-removable media.

11 In Fig. 2, the system memory 204 includes computer readable media in the
12 form of volatile, such as random access memory (RAM) 210, and/or non-volatile
13 memory, such as read only memory (ROM) 208. A basic input/output system
14 (BIOS) 212, containing the basic routines that help to transfer information
15 between elements within computer 200, such as during start-up, is stored in ROM
16 208. RAM 210 typically contains data and/or program modules that are
17 immediately accessible to and/or presently be operated on by processing unit(s)
18 202.

19 Computer 200 may further include other removable/non-removable,
20 volatile/non-volatile computer storage media. By way of example only, Fig. 2
21 illustrates a hard disk drive 228 for reading from and writing to a non-removable,
22 non-volatile magnetic media (not shown and typically called a "hard drive"), a
23 magnetic disk drive 230 for reading from and writing to a removable, non-volatile
24 magnetic disk 232 (e.g., a "floppy disk"), and an optical disk drive 234 for reading
25 from or writing to a removable, non-volatile optical disk 236 such as a CD-ROM,

1 DVD-ROM or other optical media. The hard disk drive 228, magnetic disk drive
2 230, and optical disk drive 234 are each connected to bus 206 by one or more
3 interfaces 226.

4 The drives and their associated computer-readable media provide
5 nonvolatile storage of computer readable instructions, data structures, program
6 modules, and other data for computer 200. Although the exemplary environment
7 described herein employs a hard disk 228, a removable magnetic disk 232 and a
8 removable optical disk 236, it should be appreciated by those skilled in the art that
9 other types of computer readable media which can store data that is accessible by a
10 computer, such as magnetic cassettes, flash memory cards, digital video disks,
11 random access memories (RAMs), read only memories (ROM), and the like, may
12 also be used in the exemplary operating environment.

13 A number of program modules may be stored on the hard disk 228,
14 magnetic disk 232, optical disk 236, ROM 208, or RAM 210, including, by way of
15 example, and not limitation, an operating system 214, one or more application
16 programs 216 (e.g., multimedia application program 224), other program modules
17 218, and program data 220. In accordance with the illustrated example
18 embodiment of Fig. 2, operating system 214 includes an application program
19 interface embodied as a render engine 222. As will be developed more fully
20 below, render engine 222 is exposed to higher-level applications (e.g., 216) to
21 automatically assemble filter graphs in support of user-defined development
22 projects, e.g., media processing projects. Unlike conventional media processing
23 systems, however, render engine 222 utilizes a scalable, dynamically
24 reconfigurable matrix switch to reduce filter graph complexity, thereby reducing
25 the computational and memory resources required to complete a development

1 project. Various aspects of the innovative media processing system represented by
2 a computer 200 implementing the innovative render engine 222 will be developed
3 further, below.

4 Continuing with Fig. 2, a user may enter commands and information into
5 computer 200 through input devices such as keyboard 238 and pointing device 240
6 (such as a "mouse"). Other input devices may include a audio/video input
7 device(s) 253, a microphone, joystick, game pad, satellite dish, serial port, scanner,
8 or the like (not shown). These and other input devices are connected to the
9 processing unit(s) 202 through input interface(s) 242 that is coupled to bus 206,
10 but may be connected by other interface and bus structures, such as a parallel port,
11 game port, or a universal serial bus (USB).

12 A monitor 256 or other type of display device is also connected to bus 206
13 via an interface, such as a video adapter 244. In addition to the monitor, personal
14 computers typically include other peripheral output devices (not shown), such as
15 speakers and printers, which may be connected through output peripheral interface
16 246.

17 Computer 200 may operate in a networked environment using logical
18 connections to one or more remote computers, such as a remote computer 250.
19 Remote computer 250 may include many or all of the elements and features
20 described herein relative to computer 200 including, for example, render engine
21 222 and one or more development applications 216 utilizing the resources of
22 render engine 222.

23 As shown in Fig. 2, computing system 200 is communicatively coupled to
24 remote devices (e.g., remote computer 250) through a local area network (LAN)
25 251 and a general wide area network (WAN) 252. Such networking environments

When used in a LAN networking environment, the computer 200 is connected to LAN 251 through a suitable network interface or adapter 248. When used in a WAN networking environment, the computer 200 typically includes a modem 254 or other means for establishing communications over the WAN 252. The modem 254, which may be internal or external, may be connected to the system bus 206 via the user input interface 242, or other appropriate mechanism.

In a networked environment, program modules depicted relative to the personal computer 200, or portions thereof, may be stored in a remote memory storage device. By way of example, and not limitation, Fig. 2 illustrates remote application programs 216 as residing on a memory device of remote computer 250. It will be appreciated that the network connections shown and described are exemplary and other means of establishing a communications link between the computers may be used.

Turning next to **Fig. 3**, a block diagram of an example development system architecture is presented, in accordance with one embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 3, development system 300 is shown comprising one or more application program(s) 216 coupled to render engine 222 via an appropriate communications interface 302. As used herein, application program(s) 216 are intended to represent any of a wide variety of applications which may benefit from use of render engine 222 such as, for example a media processing application 224.

The communications interface 302 is intended to represent any of a number of alternate interfaces used by operating systems to expose application program

1 interface(s) to applications. According to one example implementation, interface
2 302 is a component object model (COM) interface, as used by operating systems
3 offered by Microsoft Corporation. As introduced above, COM interface 302
4 provides a means by which the features of the render engine 222, to be described
5 more fully below, are exposed to an application program 216.

6 In accordance with the illustrated example implementation of Fig. 3, render
7 engine 222 is presented comprising source filter(s) 304A-N, transform filter(s)
8 306A-N and render filter 310, coupled together utilizing virtual pins to facilitate a
9 user-defined media processing project. According to one implementation, the
10 filters of system 300 are similar to the filters exposed in conventional media
11 processing systems. According to one implementation, however, filters are not
12 coupled via such interface pins. Rather, alternate implementations are envisioned
13 wherein individual filters (implemented as objects) make calls to other objects,
14 under the control of the render engine 222, for the desired input. Unlike
15 conventional systems, however, render engine 222 exposes a scalable, dynamically
16 reconfigurable matrix switch filter 308, automatically generated and dynamically
17 configured by render engine 222 to reduce the computational and memory
18 resource requirements often associated with development projects. As introduced
19 above, the pins (input and/or output) are application interface(s) designed to
20 communicatively couple other objects (e.g., filters).

21 In accordance with the example implementation of a media processing
22 system, an application communicates with an instance of render engine 222 when
23 the application 216 wants to process streaming media content. Render engine 222
24 selectively invokes and controls an instance of filter graph manager (not shown) to
25 automatically create a filter graph by invoking the appropriate filters (e.g., source,

transform and rendering).As introduced above, the communication of media content between filters is achieved by either (1) coupling virtual output pins of one filter to the virtual input pins of requesting filter; or (2) by scheduling object calls between appropriate filters to communicate the requested information. As shown, source filter 304 receives streaming data from the invoking application or an external source (not shown). It is to be appreciated that the streaming data can be obtained from a file on a disk, a network, a satellite feed, an Internet server, a video cassette recorder, or other source of media content. As introduced above, transform filter(s) 306 take the media content and processes it in some manner, before passing it along to render filter 310. As used herein, transform filter(s) 306 are intended to represent a wide variety of processing methods or applications that can be performed on media content. In this regard, transform filter(s) 306 may well include a splitter, a decoder, a sizing filter, a transition filter, an effects filter, and the like. The function of each of these filters is described more fully in the Griffiths application, introduced above, and generally incorporated herein by reference. The transition filter, as used herein, is utilized by render engine 222 to transition the rendered output from a first source to a second source. The effect filter is selectively invoked to introduce a particular effect (e.g., fade, wipe, audio distortion, etc.) to a media stream.

In accordance with one aspect of the embodiment, to be described more fully below, matrix switch filter 308 selectively passes media content from one or more of a scalable plurality of input(s) to a scalable plurality of output(s). Moreover, matrix switch 308 also supports implementation of a cascaded architecture utilizing feedback paths, i.e., wherein transform filters 306B, 306C, etc. coupled to the output of matrix switch 308 are dynamically coupled to one or

1 more of the scalable plurality of matrix switch input(s). An example of this
2 cascaded filter graph architecture is introduced in Fig. 3, and further explained in
3 example implementations, below.

4 Typically, media processed through source, transform and matrix switch
5 filters are ultimately passed to render filter 310, which provides the necessary
6 interface to a hardware device, or other location that accepts the renderer output
7 format, such as a memory or disk file, or a rendering device.

8 **Fig. 4** is a graphical illustration of an example software-enabled matrix
9 switch 308, according to one example embodiment of the present invention. As
10 shown, the matrix switch 308 is comprised of a scalable plurality of input(s) 402
11 and a scalable plurality of output(s) 404, wherein any one or more of the input(s)
12 402 may be iteratively coupled to any one or more of the output(s) 404, based on
13 the content of the matrix switch programming grid 406, automatically generated
14 by render engine 222. According to an alternate implementation introduced
15 above, switch matrix 308 is programmed by render engine 222 to dynamically
16 generate object calls to communicate media content between filters. In addition,
17 according to one implementation, matrix switch 308 includes a plurality of
18 input/output (I/O) buffers 408, as well as means for maintaining source, or media
19 time 410 and/or timeline, or project time 412. It is to be appreciated, however,
20 that in alternate implementations matrix switch 308 does not maintain both source
21 and project times, relying on an upstream filter to convert between these times. As
22 will be developed more fully below, matrix switch 308 dynamically couples one or
23 more of the scalable plurality of inputs 402 to one or more of the scalable plurality
24 of outputs 404 based, at least in part, on the media time 410 and/or the project time
25 412 and the content of matrix switch programming grid 406. In this regard, matrix

switch 308 may be characterized as time-aware, supporting such advanced editing features as searching/seeking to a particular point (e.g., media time) in the media content, facilitating an innovative buffering process utilizing I/O buffers 408 to facilitate look-ahead processing of media content, and the like. Thus, it will be appreciated given the discussion to follow that introduction of the matrix switch 308 provides a user with an editing flexibility that was heretofore unavailable in a personal computer-based media processing system.

As introduced above, the inputs 402 and outputs 404 of matrix switch 308 are interfaces which facilitate the time-sensitive routing of data (e.g., media content) in accordance with a user-defined development project. Matrix switch 308 has a scalable plurality of inputs 402 and outputs 404, meaning that the number of inputs 402 and outputs 404 are individually generated to satisfy a given editing project. Insofar as each of the inputs/outputs (I/O) has an associated transfer buffer (preferably shared with an adjacent filter) to communicate media content, the scalability of the input/output serves to reduce the overall buffer memory consumed by an editing project. According to one implementation, output 1 is generally reserved as a primary output, e.g., coupled to a rendering filter (not shown).

According to one implementation, for each input 402 and output 404, matrix switch 308 attempts to be the allocator, or manager of the buffer associated with the I/O(s) shared with adjacent filters. One reason is to ensure that all of the buffers are of the same size and share common attributes so that a buffer associated with any input 402 may be shared with any output 404, thereby reducing the need to copy memory contents between individual buffers associated with such inputs/outputs. If matrix switch 308 cannot be an allocator for a given

output (404), communication from an input (402) to that output is performed using a conventional memory copy operation between the individual buffers associated with the select input/output.

As introduced above, the matrix switch programming grid 406 is dynamically generated by render engine 222 based, at least in part, on the user-defined development project. As will be developed below, render engine 222 invokes an instance of filter graph manager to assemble a tree structure of an editing project, noting dependencies between source, filters and time to dynamically generate the programming grid 406. A data structure comprising an example programming grid 406 is introduced with reference to Fig. 5, below.

Turning briefly to **Fig. 5**, a graphical representation of a data structure comprising an example programming grid 406 is presented, in accordance with one embodiment of the present invention. In accordance with the illustrated example embodiment of Fig. 5, programming grid 406 is depicted as a two-dimensional data structure comprising a column along the y-axis 502 of the grid denoting input pins associated with a content chain (e.g., series of filters to process media content) of the development project. The top row along the x-axis 504 of the data structure denotes project time. With these grid “borders”, the body 506 of the grid 406 is populated with output pin assignments, denoting which input pin is coupled to which output pin during execution of the development project. In this way, render engine 222 dynamically generates and facilitates matrix switch 308. Those skilled in the art will appreciate, however, that data structures of greater or lesser complexity may well be used in support of the programming grid 406 without deviating from the spirit and scope of the present invention.

As introduced above, the filter graph is time-aware in the sense that media (source) time and project execution time are maintained. According to one implementation, matrix switch 308 maintains at least the project clock, while an upstream filter maintains the source time, converting between source and project time for all downstream filters (i.e., including the matrix switch 308). According to one implementation, the frame rate converter filter of a filter graph is responsible for converting source time to project time, and vice versa, i.e., supporting random seeks, etc. Alternatively, matrix switch 308 utilizes an integrated set of clock(s) to independently maintain project and media times.

Having introduced the architectural and operational elements of matrix switch filter 308, **Fig. 6** graphically illustrates an example filter graph implementation incorporating the innovative matrix switch 308. In accordance with the illustrated example embodiment, filter graph 600 is generated by render engine 222 in response to a user defined development project. Unlike the lengthy linear filter graphs typical of convention development systems however, filter graph 600 is shown incorporating a matrix switch filter 308 to recursively route the pre-processed content (e.g., through filters 602, 606, 610, 614 and 618, described more fully below) through a user-defined number of transform filters including, for example, transition filter(s) 620 and effects filter(s) 622. Moreover, as will be developed more fully below, the scalable nature of matrix switch filter 308 facilitates such iterative processing for any number of content threads, tracks or compositions.

According to one implementation, a matrix switch filter 308 can only process one type of media content, of the same size and at the same frame-rate (video) or modulation type/schema (audio). Thus, Fig. 6 is depicted comprising pre-processing filters with a parser filter 606 to separate, independent content type(s) (e.g., audio content and video content), wherein one of the media types would be processed along a different path including a separate instance of matrix switch 308. Thus, in accordance with the illustrated example embodiment of a media processing system, processing multimedia content including audio and video would utilize two (2) matrix switch filters 308, one dedicated to audio processing (not shown) and one dedicated to video processing. That is not to say, however, that multiple switch filters 308 could not be used (e.g., two each for audio and video) for each content type in alternate implementations. Similarly, it is anticipated that in alternate implementations a matrix switch 308 that accepts multiple media types could well be used without deviating from the spirit and scope of the present invention.

In addition filter graph 600 includes a decoder filter 610 to decode the media content. Resize filter 614 is employed when matrix switch 308 is to receive content from multiple sources, ensuring that the size of the received content is the same, regardless of the source. According to one implementation, resize filter 614 is selectively employed in video processing paths to adjust the media size of content from one or more sources to a user-defined level. Alternatively, resizer filter 614 adjusts the media size to the largest size provided by any one or more media sources. That is, if, for example, render engine 222 identifies the largest required media size (e.g., 1270x1040 video pixels per frame) and, for any content source not providing content at this size, the content is modified (e.g., stretched,

packed, etc.) to fill this size requirement. The frame rate converter (FRC) and pack filter 618, introduced above, ensures that video content from the multiple sources is arriving at the same frame rate, e.g., ten (10) frames per second. As introduced above, the FRC also maintains the distinction between source time and project time.

In accordance with one aspect of the present invention, filter graph 600 is depicted utilizing a single, negotiated buffer 604, 608, 612, 616, etc. between adjacent filters. In this regard, render engine 222 reduces the buffer memory requirements in support of a development project.

From the point of pre-processing (filters 602, 606, 610, 614, 618), rather than continue a linear filter graph incorporating all of the transition 620 and effect 622 filter(s), render engine 222 utilizes a cascade architecture, recursively passing media content through the matrix switch 308 to apply to the transform filter(s) (e.g., 620, 622, etc.) to complete the execution of the development project. It will be appreciated by those skilled in the art that the ability to recursively pass media content through one or more effect and/or transition filters provided by the matrix switch filter 308 greatly reduces the perceived complexity of otherwise large filter graphs, while reducing memory and computational overhead.

Turning to **Fig. 7**, a flow chart of an example method for generating a filter graph is presented, in accordance with one aspect of the present invention. The method 700 begins with block 702 wherein render engine 222 receives an indication to generate a filter graph representing a user-defined development project (e.g., a media editing project). According to one example implementation, the indication is received from an application 224 via COM interface(s) 302.

In block 704, render engine 222 facilitates generation of the editing project, identifying the number and type of media sources selected by the user. In block 706, based at least in part on the number and/or type of media sources, filter graph manger 222 exposes source, transform and rendering filter(s) to effect a user defined media processing project, while beginning to establish a programming grid 406 for the matrix switch filter 308.

In block 708, reflecting user editing instructions, render engine 222 completes the programming grid 406 for matrix switch 308, identifying which inputs 402 are to be coupled to which outputs 404 at particular project times.

Based, at least in part, on the programming grid 406 render engine 222 generates a matrix switch filter 308 with an appropriate number of input 402 and output 404 pins to effect the project, and assembles the filter graph, block 710.

In block 712, to reduce the buffer memory requirements for the processing project, the render engine 222 instructs the filters populating the filter graph to (re)negotiate buffer memory requirements between filters. That is, adjacent filters attempt to negotiate a size and attribute standard so that a single buffer can be utilized to couple each an output pin of one filter to an input pin of a downstream filter. An example implementation of the buffer negotiation process of block 712 is presented in greater detail with reference to Fig. 8.

Turning briefly to Fig. 8, an example method of negotiating buffer requirements between adjacent filters is presented, in accordance with one example implementation of the present invention. Once the final connection is established to matrix switch 308, matrix switch 308 identifies the maximum buffer requirements for any filter coupled to any of its pins (input 402 and/or output 404), block 802. According to one implementation, the maximum buffer requirements

In block 804, matrix switch 308 selectively removes one or more existing filter connections to adjacent filters. Matrix switch 308 then reconnects all of its pins to adjacent filters using a common buffer size between each of the pins, block 806. In block 808, matrix switch 308 negotiates to be the allocator for all of its pins (402, 404). If the matrix switch 308 cannot, for whatever reason, be the allocator for any of its input pins 402 minimal loss to performance is encountered, as the buffer associated with the input pin will still be compatible with any downstream filter (i.e., coupled to an output pin) and, thus, the buffer can still be passed to the downstream filter without requiring a memory copy operation. If, however, matrix switch 308 cannot be an allocator for one of its output pins 404, media content must then be transferred to at least the downstream filter associated with that output pin using a memory copy operation, block 810.

In block 812, once the matrix switch 308 has re-established its connection to adjacent filters, render engine 222 restores the connection in remaining filters using negotiated buffer requirements emanating from the matrix switch filter 308 buffer negotiations. Once the connections throughout the filter graph have been reconnected, the process continues with block 714 of Fig. 7.

In block 714 (Fig. 7), have re-established the connections between filters, render engine 222 is ready to implement a user's instruction to execute the media processing project.

Example Operation and Implementation(s)

The matrix switch described above is quite useful in that it allows multiple inputs to be directed to multiple outputs at any one time. These input can compete for a matrix switch output. The embodiments described below permit these competing inputs to be organized so that the inputs smoothly flow through the matrix switch to provide a desired output. And, while the inventive programming techniques are described in connection with the matrix switch as such is employed in the context of multi-media editing projects, it should be clearly understood that application of the inventive programming techniques and structures should not be so limited only to application in the field of multi-media editing projects or, for that matter, multi-media applications or data streams. Accordingly, the principles about to be discussed can be applied to other fields of endeavor in which multiple inputs can be characterized as competing for a particular output during a common time period.

In the multi-media example below, the primary output of the matrix switch is a data stream that defines an editing project that has been created by a user. Recall that this editing project can include multiple different sources that are combined in any number of different ways, and the sources that make up a project can comprise audio sources, video sources, or both. The organization of the inputs and outputs of the matrix switch are made manageable, in the examples described below, by a data structure that permits the matrix switch to be programmed.

Fig. 9 shows an overview of a process that takes a user-defined editing project and renders from it a data structure that can be used to program the matrix switch.

22 An exemplary user interface 908 is shown and represents what a user might
23 see when they produce a multimedia project with software executing on a
24 computer. In this example, the user has selected three sources A, B, and C, and
25 has assembled the sources into a project timeline. The project timeline defines

1 when the individual sources are to be rendered, as well as when any transitions
2 and/or effects are to occur.

3 In the discussion that follows, the notion of a *track* is introduced. A track
4 can contain one or more sources or source clips. If a track contains more than one
5 source clip, the source clips cannot overlap. If source clips are to overlap (e.g.
6 fading from one source to another, or having one source obscure another), then
7 multiple tracks are used. A track can thus logically represent a layer on which
8 sequential video is produced. User interface 908 illustrates a project that utilizes
9 three tracks, each of which contains a different source. In this particular project
10 source A will show for a period of time. At a defined time in the presentation,
11 source A is obscured by source B. At some later time, source B transitions to
12 source C.

13 In accordance with the described embodiment, the user-defined editing
14 project 900 is translated into a data structure 910 that represents the project. In the
15 illustrated and described example, this data structure 910 comprises a tree
16 structure. It is to be understood, however, that other data structures could be used.
17 The use of tree structures to represent editing projects is well-known and is not
18 described here in any additional detail. Once the data structure 910 is defined, it is
19 processed to provide a data structure 912 that is utilized to program the matrix
20 switch. In the illustrated and described embodiment, data structure 912 comprises
21 a grid from which the matrix switch can be programmed. It is to be understood
22 and appreciated that other data structures and techniques could, however, be used
23 to program the matrix switch without departing from the spirit and scope of the
24 claimed subject matter.

1 The processing that takes place to define data structures 910 and 912 can
2 take place using any suitable hardware, software, firmware, or combination
3 thereof. In the examples set forth below, the processing takes place utilizing
4 software in the form of a video editing software package that is executable on a
5 general purpose computer.

6 7 Example Project

8 For purposes of explanation, consider Fig. 10 which shows project 908
9 from Fig. 9 in a little additional detail. Here, a time line containing numbers 0-16
10 is provided adjacent the project to indicate when particular sources are to be seen
11 and when transitions and effects (when present) are to occur. In the examples in
12 this document, the following convention exists with respect to projects, such as
13 project 908. A priority exists for video portions of the project such that as one
14 proceeds from top to bottom, the priority increases. Thus, in the Fig. 10 example,
15 source A has the lowest priority followed by source B and source C. Thus, if there
16 is an overlap between higher and lower priority sources, the higher priority source
17 will prevail. For example, source B will obscure source A from between $t = 4-8$.

18 In this example, the following can be ascertained from the project 908 and
19 time line: from time $t=0-4$ source A should be routed to the matrix switch's
20 primary output; from $t=4-12$ source B should be routed to the matrix switch's
21 primary output; from $t=12-14$ there should be a transition between source B and
22 source C which should be routed to the matrix switch's primary output; and from
23 $t=14-16$ source C should be routed to the matrix switch's primary output. Thus,
24 relative to the matrix switch, each of the sources and the transition can be
25

characterized by where it is to be routed at any given time. Consider, for example, the table just below:

Object	Routing for a given time
C	$t = 0-12$ (nowhere); $t = 12-14$ (transition); $t = 14-16$ (primary output)
B	$t = 0-4$ (nowhere); $t = 4-12$ (primary output); $t = 12-14$ (transition); $t = 14-16$ (nowhere)
A	$t = 0-4$ (primary output); $t = 4-16$ (nowhere)
Transition	$t = 0-12$ (nowhere); $t = 12-14$ (primary output); $t = 14-16$ (nowhere)

Fig. 11 shows an exemplary matrix switch 1100 that can be utilized in the presentation of the user's project. Matrix switch 1100 comprises multiple inputs and multiple outputs. Recall that a characteristic of the matrix switch 1100 is that any of the inputs can be routed to any of the outputs at any given time. A transition element 1102 is provided and represents the transition that is to occur between sources B and C. Notice that the matrix switch includes four inputs numbered 0-3 and three outputs numbered 0-2. Inputs 0-2 correspond respectively to sources A-C, while input 3 corresponds to the output of the transition element 1102. Output 0 corresponds to the switch's primary output, while outputs 1 and 2 are routed to the transition element 1102.

The information that is contained in the table above is the information that is utilized to program the matrix switch. The discussion presented below describes but one implementation in which the information contained in the above table can be derived from the user's project time line.

Recall that as a user edits or creates a project, software that comprises a part of their editing software builds a data structure that represents the project. In the Fig. 9 overview, this was data structure 910. In addition to building the data structure that represents the editing project, the software also builds and configures a matrix switch that is to be used to define the output stream that embodies the project. Building and configuring the matrix switch can include building the appropriate graphs (e.g., a collection of software objects, or filters) that are associated with each of the sources and associating those graphs with the correct inputs of the matrix switch. In addition, building and configuring the matrix switch can also include obtaining and incorporating additional appropriate filters with the matrix switch, e.g. filters for transitions, effects, and mixing (for audio streams). This will become more apparent below.

Fig. 12 shows a graphic representation of an exemplary data structure 1200 that represents the project of Fig. 10. Here, the data structure comprises a traditional hierarchical tree structure. Any suitable data structure can, however, be utilized. The top node 1202 constitutes a *group* node. A *group* encapsulates a type of media. For example, in the present example the media type comprises video. Another media type is audio. The group node can have child nodes that are either tracks or composites. In the present example, three track nodes 1204, 1206, and 1208 are shown. Recall that each track can have one or more sources. If a track comprises more than one source, the sources cannot overlap. Here, all of the sources (A, B, and C) overlap. Hence, three different tracks are utilized for the sources. In terms of priority, the lowest priority source is placed into the tree furthest from the left at 1204a. The other sources are similarly placed. Notice that source C (1208a) has a transition 1210 associated with it. A transition object, in

1 this example, defines a two-input/one output operation. When applied to a track
2 or a composition (discussed below in more detail), the transition object will
3 operate between the track to which it has been applied, and any objects that are
4 beneath it in priority and at the same level in the tree. A “tree level” has a
5 common depth within the tree and belongs to the same parent. Accordingly, in
6 this example, the transition 1210 will operate on a source to the left of the track on
7 which source C resides, and beneath it in priority, i.e. source B. If the transition is
8 applied to any object that has nothing beneath it in the tree, it will transition from
9 blackness (and/or silence if audio is included).

10 Once a data structure representing the project has been built, in this case a
11 hierarchical tree structure, a rendering engine processes the data structure to
12 provide another data structure that is utilized to program the matrix switch. In the
13 Fig. 9 example, this additional data structure is represented at 912. It will be
14 appreciated and understood that the nodes of tree 1200 can include so-called meta
15 information such as a name, ID, and a time value that represents when that
16 particular node’s object desires to be routed to the output, e.g. node 1204a would
17 include an identifier for the node associating it with source A, as well as a time
18 value that indicates that source A desires to be routed to the output from time $t = 0$ -
19 8. This meta information is utilized to build the data structure that is, in turn,
20 utilized to program the matrix switch.

21 In the example about to be described below, a specific data structure in the
22 form of a grid is utilized. In addition, certain specifics are described with respect
23 to how the grid is processed so that the matrix switch can be programmed. It is to
24 be understood that the specific described approach is for exemplary purposes only
25 and is not intended to limit application of the claims. Rather, the specific approach

1 constitutes but one way of implementing broader conceptual notions embodied by
2 the inventive subject matter.

3 Figs. 13-18 represent a process through which the inventive grid is built. In
4 the grid about to be described, the x axis represents time, and the y axis represents
5 layers in terms of priority that go from lowest (at the top of the grid) to highest (at
6 the bottom of the grid). Every row in the grid represents the video layer.
7 Additionally, entries made within the grid represent output pins of the matrix
8 switch. This will become apparent below.

9 The way that the grid is built in this example is that the rendering engine
10 does a traversal operation on the tree 1200. In this particular example, the
11 traversal operation is known as a “depth-first, left-to-right” traversal. This
12 operation will layerize the nodes so that the leftmost track or source has the lowest
13 priority and so on. Doing the above-mentioned traversal on tree 1200 (Fig. 12),
14 the first node encountered is node 1204 which is associated with source A. This is
15 the lowest priority track or source. A first row is defined for the grid and is
16 associated with source A. After the first grid row is defined, a grid entry is made
17 and represents the time period for which source A desires to be routed to the
18 matrix switch’s primary output.

19 Fig. 13 shows the state of a grid 1300 after this first processing step.
20 Notice that from time $t = 0-8$, a “0” has been placed in the grid. The “0”
21 represents the output pin of the matrix switch—in this case the primary output.
22 Next, the traversal encounters node 1206 (Fig. 12) which is associated with source
23 B. A second row is thus defined for the grid and is associated with source B.
24 After the second grid row is defined, a grid entry is made and represents the time
25

period for which source B desires to be routed to the matrix switch's primary output.

Fig. 14 shows the state of grid 1300 after this second processing step. Notice that from time $t = 4-14$, a "0" has been placed in the grid. Notice at this point that something interesting has occurred which will be resolved below. Each of the layers has a common period of time (i.e. $t = 4-8$) for which it desires to be routed to the matrix switch's primary output. However, because of the nature of the matrix switch, only one input can be routed to the primary output at a time. Next, the traversal encounters node 1208 (Fig. 12) which is associated with source C. In this particular processing example, a rule is defined that sources on tracks are processed before transitions on the tracks are processed because transitions operate on two objects that are beneath them. A third row is thus defined for the grid and is associated with source C. After the third row is defined, a grid entry is made and represents the time period for which source C desires to be routed to the matrix switch's primary output.

Fig. 15 shows the state of grid 1300 after this third processing step. Notice that from time $t = 12-16$, a "0" has been placed in the grid. Next, the traversal encounters node 1210 (Fig. 12) which corresponds to the transition. Thus, a fourth row is defined in the grid and is associated with the transition. After the fourth row is defined, a grid entry is made and represents the time period for which the transition desires to be routed to the matrix switch's primary output.

Fig. 16 shows the state of grid 1300 after this fourth processing step. Notice that from time $t = 12-14$, a "0" has been placed in the grid for the transition entry. The transition is a special grid entry. Recall that the transition is programmed to operate on two inputs and provide a single output. Accordingly,

Fig. 17 shows the state of the grid at this point in the processing. Next, a pruning function is implemented which removes any other lower priority entry that is contending for the output with a higher priority entry. In the example, the portion of A from $t=4-8$ gets removed because the higher priority B wants the output for that time.

Fig. 18 shows the grid with a cross-hatched area that signifies that portion of A's grid entry that has been removed.

As alluded to above, one of the innovative aspects of the matrix switch 308 is its ability to seek to any point in a source, without having to process the intervening content serially through the filter. Rather, matrix switch 308 identifies an appropriate transition point and dumps at least a subset of the intervening content, and continues processing from the sought point in the content.

The ability of the matrix switch 308 to seek to any point in the media content gives rise to certain performance enhancement heretofore unavailable in

computer implemented media processing systems. For example, generation of a filter graph by render engine 222 may take into account certain performance characteristics of the media processing system which will execute the user-defined media processing project. In accordance with this example implementation, render engine 222 may access and analyze the system registry of the operating system, for example, to ascertain the performance characteristics of hardware and/or software elements of the computing system implementing the media processing system, and adjust the filter graph construction to improve the perceived performance of the media processing system by the user. Nonetheless, there will always be a chance that a particular instance of a filter graph will not be able to process the media stream fast enough to provide the desired output at the desired time, i.e., processing of the media stream bogs down leading to delays at the rendering filter. In such a case, matrix switch 308 will recognize that it is not receiving media content at the appropriate project time, and may skip certain sections of the project in an effort to “catch-up” and continue the remainder of the project in real time. According to one implementation, when matrix switch 308 detects such a lag in processing, it will analyze the degree of the lag and issue a seek command to the source (through the source processing chain) to a future point in the project, where processing continues without processing any further content prior to the sought point.

Thus, for the editing project depicted in Fig. 10, the processing described above first builds a data structure (i.e. data structure 1200 in Fig. 12) that represents the project in hierarchical space, and then uses this data structure to define or create another data structure that can be utilized to program the matrix switch.

Fig. 19 is a flow diagram that describes steps in a method in accordance with the described embodiment. The method can be implemented in any suitable hardware, software, firmware, or combination thereof. In the illustrated and described embodiment, the method is implemented in software.

Step 1900 provides a matrix switch. An exemplary matrix switch is described above. Step 1902 defines a first data structure that represents the editing project. Any suitable data structure can be used, as will be apparent to those of skill in the art. In the illustrated and described embodiment, the data structure comprises a hierarchical tree structure having nodes that can represent tracks (having one or more sources), composites, transitions and effects. Step 1904 processes the first data structure to provide a second data structure that is configured to program the matrix switch. Any suitable data structure can be utilized to implement the second data structure. In the illustrated and described embodiment, a grid structure is utilized. Exemplary processing techniques for processing the first data structure to provide the second data structure are described above. Step 1906 then uses the second data structure to program the matrix switch.

Example Project with a Transition and an Effect

Consider project 2000 depicted in Fig. 20. In this project there are three tracks, each of which contains a source, i.e. source A, B and C. This project includes an effect applied on source B and a transition between sources B and C. The times are indicated as shown.

As the user creates their project, a data structure representing the project is built. Fig. 21 shows an exemplary data structure in the form of a hierarchical tree

2100 that represents project 2000. There, the data structure includes three tracks, each of which contains one of the sources. The sources are arranged in the tree structure in the order of their priority, starting with the lowest priority source on the left and proceeding to the right. There is an effect (i.e. "Fx") that is attached to or otherwise associated with source B. Additionally, there is a transition attached to or otherwise associated with source C.

In building the grid for project 2000, the following rule is employed for effects. An effect, in this example, is a one-input/one-output object that is applied to one object—in this case source B. When the effect is inserted into the grid, it looks for any one object beneath it in priority that has a desire to be routed to the primary output of the matrix switch at the same time. When it finds a suitable object, it redirects that object's output from the matrix switch's primary output to an output associated with the effect.

As an example, consider Fig. 22 and the grid 2200. At this point in the processing of tree 2100, the rendering engine has incorporated entries in the grid corresponding to sources A, B and the effect. It has done so by traversing the tree 2100 in the above-described way. In this example, the effect has already looked for an object beneath it in priority that is competing for the primary output of the matrix switch. It found an entry for source B and then redirected B's grid entry to a matrix switch output pin that corresponds to the effect—here output pin 1.

As the render engine 222 completes its traversal of tree 2100, it completes the grid. Fig. 23 shows a completed grid 2200. Processing of the grid after that which is indicated in Fig. 22 takes place substantially as described above with respect to the first example. Summarizing, this processing though: after the effect is entered into the grid and processed as described above, the traversal of tree 2100

next encounters the node associated with source C. Thus, a row is added in the grid for source C and an entry is made to indicate that source C desires the output from $t = 12-16$. Next, the tree traversal encounters the node associated with the transition. Accordingly, a row is added to the grid for the transition and a grid entry is made to indicate that the transition desires the output from $t = 12-14$. Now, as described above, the grid is examined to find two entries, lower in priority than the transition and located at the same tree level as the transition, that compete for the primary output of the matrix switch. Here, those entries correspond to the grid entries for the effect and source C that occur from $t = 12-14$. These grid entries are thus redirected to output pins of the matrix switch 308 that correspond to the transition—here pins 2 and 3 as indicated. Next, the grid is pruned which, in this example, removes a portion of the grid entry corresponding to source A for $t = 4-8$ because of a conflict with the higher-priority entry for source B.

Fig. 24 shows the resultant matrix switch that has been built and configured as the grid was being processed above. At this point, the grid can be used to program the matrix switch. From the grid picture, it is very easy to see how the matrix switch 308 is going to be programmed. Source A will be routed to the matrix switch's primary output (pin 0) from $t = 0-4$; source B will be redirected to output pin 1 (effect) from $t = 4-14$ and the effect on B will be routed to the output pin 0 from $t = 4-12$. From $t = 12-14$, the effect and source C will be routed to output pins corresponding to the transition (pins 2 and 3) and, accordingly, during this time the transition (input pin 4) will be routed to the primary output (output pin 0) of the matrix switch. From $t = 14-16$, source C will be routed to the primary output of the matrix switch.

It will be appreciated that as the software, in this case the render engine 222, traverses the tree structure that represents a project, it also builds the appropriate graphs and adds the appropriate filters and graphs to the matrix switch. Thus, for example, as the render engine 222 encounters a tree node associated with source A, in addition to adding an entry to the appropriate grid, the software builds the appropriate graphs (i.e. collection of linked filters), and associates those filters with an input of the matrix switch. Similarly, when the render engine 222 encounters an effect node in the tree, the software obtains an effect object or filter and associates it with the appropriate output of the matrix switch. Thus, in the above examples, traversal of the tree structure representing the project also enables the software to construct the appropriate graphs and obtain the appropriate objects and associate those items with the appropriate inputs/outputs of the matrix switch 308. Upon completion of the tree traversal and processing of the grid, an appropriate matrix switch has been constructed, and the programming (i.e. timing) of inputs to outputs for the matrix switch has been completed.

Treatment of “blanks” in a Project

There may be instances in a project when a user leaves a blank in the project time line. During this blank period, no video or audio is scheduled for play.

Fig. 25 shows a project that has such a blank incorporated therein. If there is such a blank left in a project, the software is configured to obtain a “black” source and associate the source with the matrix switch at the appropriate input pin. The grid is then configured when it is built to route the black source to the output at the appropriate times and fade from the black (and silent) source to the next

1 source at the appropriate times. The black source can also be used if there is a
2 transition placed on a source for which there is no additional source from which to
3 transition.

4 5 **Audio Mixing**

6 In the examples discussed above, sources comprising video streams were
7 discussed. In those examples, at any one time, only two video streams were
8 combined into one video stream. However, each project can, and usually does
9 contain an audio component. Alternately, a project can contain only an audio
10 component. The audio component can typically comprise a number of different
11 audio streams that are combined. The discussion below sets forth but one way of
12 processing and combining audio streams.

13 In the illustrated example, there is no limit on the number of audio streams
14 that can be combined at any one time.

15 Suppose, for example, there is an audio project that comprises 5 tracks, A-
16 E. Fig. 26 shows an exemplary project. The shaded portions of each track
17 represent the time during which the track is not playing. So, for example, at $t=0-4$,
18 tracks B, D, and E are mixed together and will play. From $t = 4-10$, tracks A-E are
19 mixed together and will play, and the like.

20 Fig. 27 shows the grid for this project at 2700. Since we are dealing with
21 this composition now, all of the effects and transitions including the audio mixing
22 are only allowed to affect things in this composition. Thus, there is the concept of
23 a boundary 2702 that prevents any actions or operations in this composition from
24 affecting any other grid entries. Note that there are other entries in the grid and
25

time period in the grid, each of the sources is routed to the mix element which, in turn, mixes the appropriate audio streams and presents them to the primary output pin 0 of the matrix switch.

Compositions

There are situations that can arise when building an editing project where it would be desirable to apply an effect or a transition on just a subset of a particular project or track. Yet, there is no practicable way to incorporate the desired effect or transition. In the past, attempts to provide added flexibility for editing projects have been made in the form of so called “bounce tracks”, as will be appreciated and understood by those of skill in the art. The use of bounce tracks essentially involves processing various video layers (i.e. tracks), writing or moving the processed layers or tracks to another location, and retrieving the processed layers when later needed for additional processing with other layers or tracks. This type of processing can be slow and inefficient.

To provide added flexibility and efficiency for multi-media editing projects, the notion of a *composite or composition* is introduced. A composite or composition can be considered as a representation of an editing project as a single track. Recall that editing projects can have one or more tracks, and each track can be associated with one or more sources that can have effects applied on them or transitions between them. In addition, compositions can be nested inside one another.

Example Project with Composite

Consider, for example, Fig. 29 which illustrates an exemplary project 2900 having a composition 2902. In this example, composition 2902 comprises sources B and C and a transition between B and C that occurs between $t = 12-14$. This composition is treated as an individual track or layer. Project 2900 also includes a source A, and a transition between source A and composition 2902 at $t = 4-8$. It will be appreciated that compositions can be much more complicated than the illustrated composition, which is provided for exemplary purposes only. Compositions are useful because they allow the grouping of a particular set of operations on one or more tracks. The operation set is performed on the grouping, and does not affect tracks that are not within the grouping. To draw an analogy, a composition is similar in principle to a mathematical parenthesis. Those operations that appear within the parenthesis are carried out in conjunction with those operations that are intended to operate of the subject matter of the parenthesis. The operations within the parenthesis do not affect tracks that do not appear within the parenthesis.

In accordance with the processing that is described above in connection with Fig. 19, a first data structure is defined that represents the editing project. Fig. 30 shows an exemplary data structure 3000 in the form of a hierarchical tree structure. In this example, group node 3002 includes two children—track node 3004 and composite node 3006. Track node 3004 is associated with source A. Composite node 3006 includes two children—track nodes 3008 and 3010 that are respectively associated with sources B (3008a) and C (3010a). A transition T2 (3012) is applied on source C and a transition T1 (3014) is applied on composition 3006.

1 Next, data structure 3000 is processed to provide a second data structure
2 that is configured to program the matrix switch. Note that as the data structure is
3 being programmed, a matrix switch is being built and configured at the same time.
4 In this example, the second data structure comprises a grid structure that is
5 assembled in much the same way as was described above. There are, however,
6 some differences and, for purposes of understanding, the complete evolution of the
7 grid structure is described here. In the discussion that follows, the completed
8 matrix switch is shown in Fig. 38.

9 When the rendering engine initiates the depth-first, left-to-right traversal of
10 data structure 3000, the first node it encounters is track node 3004 which is
11 associated with source A. Thus, a first row of the grid is defined and a grid entry
12 is made that represents the time period for which source A desires to be routed to
13 the matrix switch's primary output pin.

14 Fig. 31 shows the state of a grid 3100 after this first processing step. Next
15 the traversal of data structure 3000 encounters the composite node 3006. The
16 composite node is associated with two tracks—track 3008 and track 3010. Track
17 3008 is associated with source B. Accordingly, a second row of the grid is defined
18 and a grid entry is made that represents the time period for which source B desires
19 to be routed to the matrix switch's primary output pin. Additionally, since B is a
20 member of a composition, meta-information is contained in the grid that indicates
21 that this grid row defines one boundary of the composition. This meta-
22 information is graphically depicted with a bracket that appears to the left of the
23 grid row.

24 Fig. 32 shows the state of grid 3100 after this processing step. Next, the
25 traversal of data structure 3000 encounters node 3010 which is associated with

1 source C. Thus, a third row of the grid is added and a grid entry is made that
2 represents the time period for which source C desires to be routed to the matrix
3 switch's primary output pin.

4 Fig. 33 shows the state of grid 3100 after this processing step. Notice that
5 the bracket designating the composition now encompasses the grid row associated
6 with source C. The traversal next encounters node 3012 which is the node
7 associated with the *second* transition T2. Thus, as in the above example, a grid
8 row is added for the transition and a grid entry is made that represents the time
9 period for which the transition desires to be routed to the matrix switch's primary
10 output pin.

11 Fig. 34 shows the state of grid 3100 after this processing step. Notice that
12 the bracket designating the composition is now completed and encompasses grid
13 row entries that correspond to sources B and C and the transition between them.
14 Recall from the examples above that a transition, in this example, is programmed
15 to operate on two inputs and provide a single output. In this instance, and because
16 the transition occurs within a composition, the transition is constrained by a rule
17 that does not allow it to operate on any elements outside of the composition.
18 Thus, starting at the transition entry and working backward through the grid,
19 entries at the same tree level and within the composition (as designated by the
20 bracket) are examined to ascertain whether they contain entries that indicate that
21 they want to be routed to the output during the same time that the transition is to
22 be routed to the output. Here, both of the entries for sources B and C have
23 portions that conflict with the transition's entry. Accordingly, those portions of
24 the grid entries for sources B and C are redirected or changed to correspond to
25

1 output pins that are associated with a transition element that corresponds to
2 transition T2.

3 Fig. 35 shows the state of grid 3100 after this processing step. The
4 traversal next encounters node 3014 which is the node that is associated with the
5 transition that occurs between source A and composition 2902 (Fig. 29).
6 Processing of this transition is similar to processing of the transition immediately
7 above except for the fact that the transition does not occur within the composition.
8 Because the transition occurs between the composition and another source, one of
9 the inputs for the transition will be the composition, and one of the inputs will be
10 source A (which is outside of the composition). Thus, a grid row is added for this
11 transition and a grid entry is made that represents the time period for which the
12 transition desires to be routed to the matrix switch's primary output pin.

13 Fig. 36 shows the state of grid 3100 after this processing step. At this point
14 then, the grid is examined for entries that conflict with the entry for transition T1.
15 One conflicting grid entry is found for the row that corresponds to source B (inside
16 the composition) and one that corresponds to source A (outside the composition).
17 Accordingly, those portions of the grid row that conflict with transition T1 are
18 changed or redirected to have values that are associated with output pins of the
19 matrix switch that are themselves associated with a transition element T1. In this
20 example, redirection causes an entry of "3" and "4" to be inserted as shown.

21 Fig. 37 shows the state of grid 3100 after this processing step. If necessary,
22 a pruning operation would further ensure that the grid has no competing entries for
23 the primary output of the matrix switch. The associated input pin numbers of the
24 matrix switch are shown to the left of grid 3100.

1 Fig. 38 shows a suitably configured matrix switch that has been build in
2 accordance with the processing described above. Recall that, as data structure
3 3000 (Fig. 30) is processed by the rendering engine, a matrix switch is built and
4 configured in parallel with the building and processing of the grid structure that is
5 utilized to program the matrix switch. From the matrix switch and grid 3100 of
6 Fig. 37, the programming of the switch can be easily ascertained.

7 Fig. 38a shows an exemplary data structure that represents a project that
8 illustrates the usefulness of composites. In this example, the project can
9 mathematically be represented as follows:

10
11 (Fx-noisy (A Tx-Blend B)) Tx-Blend C
12

13 Here, an effect (noisy) is applied to A blended with B, the result of which is
14 applied to a blend with C. The composite in this example allows the grouping of
15 the things beneath it so that the effect (noisy), when it is applied, is applied to
16 everything that is beneath it. Notice that without the composite node, there is no
17 node where an effect can be applied that will affect (A Tx-Blend B). Hence, in
18 this example, operations that appear within the parenthesis are carried out on
19 tracks that appear within the parenthesis. Those operations do not affect tracks
20 that are not within the parenthesis.

21 Fig. 39 is a flow diagram that described steps in a method in accordance
22 with one embodiment. The method can be implemented in any suitable hardware,
23 software, firmware, or combination thereof. In the presently-described example,
24 the method is implemented in software.
25

Step 3900 defines a multimedia editing project that includes at least one composite. The composite represents multiple tracks as a single track for purposes of the processing described just below. It is important to note that, in the processing described just below, and because of the use of composites, the extra processing that is required by bounce tracks is avoided (i.e. operating on two tracks, moving the operation result to another location, and retrieving the operation result when later needed). This reduces the processing time that is required to render a multi-media project. Step 3902 defines a first data structure that represents the editing project. Any suitable data structure can be utilized. In the present example, a data structure in the form of a hierarchical tree is utilized. An exemplary tree is shown in Fig. 30. Step 3904 processes the first data structure to provide a second data structure that is configured to program a matrix switch. In the illustrated example, the second data structure comprises a grid structure. Exemplary processing is described in the context of Figs. 30-37. Step 3906 then programs the matrix switch using the second data structure.

Implementing Dynamic Properties on Objects that Only Support Static Properties

Computer-implemented objects typically have properties associated with them. Many different types of properties can be associated with such objects. For example, consider a computer-implemented object in the form of a ball. The ball can have various properties such as color, direction of movement, speed, etc. These properties typically have values associated with them. For example, for the property "color", property values might include blue, orange, red, and yellow. For the property "direction", property values might include x-direction, y-direction, and z-direction.

Many objects have so-called static properties whose values can only be set at a current time. In essence, these types of objects have to be notified to change their property values currently. These types of objects are not pre-programmable to change their property values at some point in the future. Dynamic properties are properties whose values can change at certain times. An object such as a ball that supports dynamic properties can, for example, be pre-programmed to change its color, direction, speed, and the like, dynamically at any time in the future. That is, such objects can be preprogrammed to change their property values.

It would be desirable to be able to impart to objects that support only static properties characteristics such that they behave like they support dynamic properties. That is, it would be desirable to be able to pre-program objects that support only static properties.

Fig. 40 is a block diagram of a system that can be utilized, in accordance with the described embodiment, to impart to objects that support only static properties, characteristics that enable them to behave in a way that appears as if they support dynamic properties.

In the illustrated example, an object 4000 is provided and constitutes any object that can support only static properties. Accordingly, a set of static properties 4002 is shown. It will be appreciated that any suitable number of objects such as object 4000 can be provided. In addition, any number and/or type of suitable static properties can be supported by the object. For purposes of this example, however, assume that object 4000 represents a ball and that the ball has properties that include color, directions of movement and speed.

A helper object 4004 is provided and is used in connection with object 4000 to, in a sense, imitate or simulate dynamic properties. The result of the interaction

Thus, in the described embodiment, an object that does not support dynamic properties is made to appear as if it supports dynamic properties through the use of and interaction with a helper object. The helper object can maintain various property data that it then uses to effect a property value change on its associated object. The property data can include, without limitation, the time when the change is to take effect, type of property value change, and how to effect the change. In this manner, property values and changes in those values can be pre-programmed.

In the illustrated and described embodiment, helper object 4004 makes use of data structures that assist it in its job. For each property whose values are to be changed, a set of structures is defined. Each structure comprises entries for time (i.e. the time that a property value change is to take place), a variant for the value of the property, and an entry for how the change is to take place (e.g. immediately, or in some interpolated fashion). Thus, in the Fig. 40 example, structure array 4006 includes four sets of structures designated set 1, set 2, set 3, and set 4. Each set is associated with a different property whose value can be changed on an object that supports only static properties. For example, one set each of structures might be associated with an object's speed, color, and x/y direction of movement.

Consider the following example where object 4000 comprises a ball. Consider that, for the ball, we want the helper object to cause the ball to be red at time 0, change immediately to blue at time 2, and blend into yellow at time 4. An exemplary data structure that can enable this to be done is shown in Fig. 40. In this example, the data structure “(0, Red, Jump)” indicates that at time 0, the ball is to immediately (i.e. “jump to”) be red; the structure “(2, Blue, Jump)” indicates that at time 2, the ball is to immediately be blue; and the structure “(4, Yellow, Interpolate)” indicates that at time 4, the ball is to change in a linear fashion (i.e. interpolate) into yellow from blue.

The helper object 4004 is programmed (or pre-programmed) with the various data structures that define the desired properties, their value changes, and the time and manner in which the value change is to take place. The helper object keeps track of the time, by for example, being called or otherwise notified of the time. At the appropriate times, the helper object, using the array of structures, calls the object whose property values are desired to be changed so that the property values can be changed. In the present example, at time 0, the helper object calls object 4000 or a property object on the object itself, and causes the property value for the color property to be set to “red”. At time 2, the helper object 4004 calls the object 4000 and causes the property value for the color property to be set to “blue”. At time 4, the helper object calls the object 4000 and causes the property value for the color property to interpolate from its previous value (i.e. blue) to the present value, i.e. yellow.

Thus, object 4000 appears as if it supports dynamic properties when, in fact, it only supports static properties.

As another example, consider Fig. 41. In this example a ball 4100 has directional properties that allows it to move within a coordinate system in both the x and y directions. The ball constitutes an object that supports only static properties, i.e. it cannot be pre-programmed for movement in the future. The coordinates range from 0-200 in both the x and y directions as shown. Each property (movement in the y-direction and movement in the x-direction) is associated with a set of data structures as indicated generally at 4102. In the present example, at time 0, the ball is located at x=0, y=100; at time 2, the ball jumps to x=100 and maintains its y position at 100. At time 4, the ball's position is interpolated (linearly) from x=100 to x=200, and from y=100 to y=200.

Fig. 42 is a flow diagram that describes steps in a method in accordance with the described embodiment. The method can be implemented in any suitable hardware, software, firmware or combination thereof. In the illustrated example, certain steps are implemented in software. The illustrated method can relieve a top level application from the burden of having to set object properties over and over again by keeping a thread running. The top level application relies on the helper object to know all of the properties that need to be changed or set and the time at which to change or set the properties. The helper object will then do the work involved with changing or setting the properties. This includes having a thread running and continuously setting or changing the properties as instructed by the application. In performing its job, the helper object can access a clock that the application is using so that it knows what time it is so that it can, accordingly, set or change the right property.

Step 4200 provides one or more objects that support only static properties. These objects can be provided by an application executing on a suitable computer

processor. Step 4202 provides at least one programmable helper object. The helper object can be provided by the application program. The helper object is desirably configured to be used to effect property value changes on the objects that support only static properties. Step 4204 programs the helper object with property data that is to be used by the helper object to effect one or more property value changes on one or more of the objects that support only static properties. This step can be implemented by an application program accepting input from a user that defines the property data. Specific examples of property data are given above and can include, without limitation, time when the property value change is to take effect, type of property value change, and how to effect the property value change. The property data can be organized in any suitable way. For example, in the discussion above, an array of data structures is used to hold the property data. Each structure array has one or more sets of data structures, and each set of data structures is associated with a property that is to be changed. It is to be appreciated, however, that any suitable method, structure, or technique of effecting the property changes can be used. Step 4206 then calls, with the helper object(s), one or more of the objects that support only static properties to effect the property value change on the object(s). In the illustrated and described embodiment, the helper object is called with the current time so that it knows when to call the objects whose property values it desires to change.

Specific Example

The inventive helper object can be used in many contexts where it is desirable to impart to objects that support only static properties, characteristics or behavior that appears dynamic in nature. The helper object can relieve the human

1 programmer of a great burden which, in the past, has required programmers to
2 specifically provide the software code (including timing) to implement all of the
3 property value changes. One particularly useful context for application of the
4 above-described techniques is in the area of multi-media and, in particular, the
5 editing of multi-media projects.

6 Fig. 43 shows a matrix switch 4300 that has been built as described above
7 and, as shown, is configured to effect a transition between two multi-media
8 streams. To perform that task, a transition object 4302 is provided. The transition
9 object is shown to include a helper object 4304 having a structure array, and an
10 object 4306 that supports only static properties. Object 4306 is the actual
11 transform that performs the desired transition and, in this example, is implemented
12 as a DirectX transform, although any suitable transform can be used.

13 In this example, transition object 4302 has a property called "percent
14 complete" and represents how complete the transition is for a given project time.
15 There are different percentages (i.e. values) of completion that vary as the project
16 time progresses. For example, at time 1, the transition may be 1% complete, at
17 time 2 the transition may be 3% complete and the like. Object 4306 supports only
18 static properties and, hence, is called by the helper object 4304 during the course
19 of a transition and told, by the helper object, to render a transition at a particular
20 percentage of completion. The helper object 4304 is programmed by data that is
21 received from one or more of the matrix switch, the grid and filter graphs that are
22 assembled to complete the project. This information is in a form that is not
23 understood by the object 4306 that performs the actual transition. Accordingly, the
24 helper object 4304 acts as an intermediary in that it receives programming
25 information that describes when, how, and what properties and values are to be

1 changed. The helper object then translates this programming information into
2 instructions that are specific to and understood by object 4306 that performs the
3 actual transition.

4 Accordingly, in this example, helper object 4304 defines an intermediate
5 level that provides the ability to simulate pre-programmed dynamic properties on
6 an object (i.e. object 4306) that otherwise does not have the capability to have
7 dynamic properties.

8 9 **Conclusion**

10 The described embodiments can be used to provide improvements over
11 previous multi-media editing systems. Various efficiencies are achieved that
12 reduce the processing times and can thereby improve the user experience when
13 using multi-media project editing software applications.

14 Although the invention has been described in language specific to structural
15 features and/or methodological steps, it is to be understood that the invention
16 defined in the appended claims is not necessarily limited to the specific features or
17 steps described. Rather, the specific features and steps are disclosed as preferred
18 forms of implementing the claimed invention.